DOCUMENT RESUME

ED 217 407 CS 006 730

AUTHOR Malt, Barbara C.; Smith, Edward E.
TITLE The Role of Familiarity in Determining Typicality. Technical Report No. 250.
INSTITUTION Bolt, Beranek and Newman, Inc., Cambridge, Mass.; Illinois Univ., Urbana. Center for the Study of Reading.
SPONS AGENCY National Inst. of Education (ED), Washington, DC.; Public Health Service (DHEW), Rockville, Md.
PUB DATE Jun 82
CONTRACT 400-76-0116
GRANT PHS-MH-19705
NOTE 29p.

EDRS PRICE MF01/PC02 Plus Postage.
DESCRIPTORS *Classification; *Cognitive Processes; Higher Education; *Measurement Techniques; *Reading Research; Research Methodology; *Word Recognition
IDENTIFIERS *Familiarity

ABSTRACT

M.H. Ashcraft found that people tend to know more properties of items they rate as typical of a category than of items they rate as atypical, suggesting that variations in typicality result from variations in familiarity. Three experiments were designed to challenge this suggestion. The first investigated whether familiarity is necessarily correlated with typicality ratings for a large sample of category members that span the typicality range. In this experiment, 20 college students generated properties for category members and 19 other students rated the category members for typicality. Approximately 300 students participated in the second experiment, which tested whether, for a random sample of 15 items from each of 8 categories, a positive correlation would be found between typicality ratings and number of properties listed. The third experiment produced a set of typicality ratings in which the 20 subjects were able to indicate that they were unable to rate the item because they were unfamiliar with it. Results showed that (1) subjects sometimes produced more properties for items they rated low in typicality; (2) in a large, random sample of items, subjects tended to produce fewer properties for atypical items; and (3) subjects tended to assign totally unfamiliar words to the bottom of a typicality scale rather than reflect low typicality of the referents themselves. (FL)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED217407

U.S. DEPARTMENT OF EDUCATION
NATIONAL INSTITUTE OF EDUCATION
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)
X This document has been reproduced as received from the person or organization originating it.
☐ Minor changes have been made to improve reproduction quality.
• Points of view or opinions stated in this document do not necessarily represent official NIE position or policy.

CENTER FOR THE STUDY OF READING

Technical Report No. 250

THE ROLE OF FAMILIARITY IN DETERMINING TYPICALITY

Barbara C. Malt
Stanford University

Edward E. Smith
Bolt Beranek and Newman Inc.

June 1982

University of Illinois
at Urbana-Champaign
51 Gerty Drive
Champaign, Illinois 61820

Bolt Beranek and Newman Inc.
50 Moulton Street
Cambridge, Massachusetts 02238

The research reported herein was supported in part by the National Institute of Education under Contract No. HEW-NIE-C-400-76-0116, and in part by U.S. Public Health Service Grant MH-19705. This paper will appear in the journal Memory and Cognition in 1981.

CS006730

EDITORIAL BOARD

William Nagy and Stephen Wilhite
Co-Editors

Harry Blanchard
Charlotte Blomeyer
Nancy Bryant
Larry Colker
Avon Crismore
Meg Gallagher
Anne Hay
Asghar Iran-Nejad
Margi Laff
Jill LaZansky
Cindy Steinberg
Terry Turner
Janet Williams
Paul Wilson

Abstract

Ashcraft (1978b) found that people tend to know more properties of instances they rate as typical of a category than ot instances they rate as atypical. This suggests that variations in typicality result from variations in familiarity. We present three experiments that challenge or qualify this suggestion. Experiment 1 showed that subjects sometimes produce more properties for items they rate low in typicality. Experiment 2 showed that in a large, random sample of items, there was a tendency to produce fewer properties for atypical items, but Experiment 3 indicated that part of the reason for this result was a response bias to assign totally unfamiliar words to the bottom of the typicality scale, rather than reflecting low typicality of the referents themselves.

The Role of Familiarity in Determining Typicality

All members of a semantic category are not equally representative or typical of that category: A peach is a more typical fruit than a pomegranate, and a robin a more typical bird than a roadrunner. By now it is well-established that people show strong agreement in their ratings of how typical members of semantic categories are, and that rated typicality predicts performance in a wide variety of tasks such as reaction time to verify category membership, order of learning by children, and order and probability of output in free listing to category names (Mervis & Rosch, in press; Rips, Shoben, & Smith, 1973; Rosch, 1975; Rosch & Mervis, 1975). Understanding what determines typicality is thus a step toward understanding the principles by which information in semantic memory is acquired and organized.

Most explanations of typicality have focused on the internal structure of categories, specifically on the properties of the category members and/or those of the category itself (e.g., Rosch & Mervis, 1975; Smith, Shoben, & Rips, 1974). Rosch and Mervis (1975), for example, suggested that typicality is based on the distribution of properties among category members, where typical members have many properties that occur frequently across category members, and atypical members have properties that are less

frequent among category members. In support of this, Rosch and Mervis found a strong correlation between typicality ratings and family resemblance scores, a measure of overlap of an exemplar's properties with the properties of other category members.

An alternative to a structural account is a familiarity-based explanation of typicality. For most college students, peaches are more commonly encountered than pomegranates, and robins more commonly encountered than roadrunners, and those category members that are most frequently seen, talked about, or interacted with, will be those judged most typical. Until recently, this kind of explanation has been discounted, largely on the basis of data on word frequency and artificial category experiments. Mervis, Catlin, and Rosch (1976), for example, found no correlation between Kucera and Francis (1967) word frequency and rated typicality for members of common categories. And Rosch, Simpson, and Miller (1976) demonstrated that when frequency of presentation was manipulated for members of artificial categories, typicality ratings reflected property relations rather than frequency. None of these results, however, directly addressed the problem of what determines typicality for real-world categories. Frequency counts of written prose need not necessarily reflect how familiar or common in the environment a category member is; we may tend, for example, to write disproportionately about the uncommon

and unusual. It also is not clear to what extent the artificial category results are indicative of the real-world phenomenon, since there may be a much greater range in familiarity for members of real-world categories than was used in the artificial category research.

Recent work, however, has revived familiarity explanations of typicality. McCloskey (1980) had subjects rate the familiarity of meaning of words used in several published semantic-memory experiments; he found a wide range of familiarity although word frequency had been carefully controlled in these experiments. This familiarity measure seemed to correlate with rated typicality, and the effect of typicality on reaction time was reduced substantially (though not eliminated) when familiarity was partialed out.

It is not clear in McCloskey's study whether subjects were rating how familiar the word was or how familiar the real-world referent was. Ashcraft (1978b) investigated familiarity of referents of words more directly. He had subjects free-list properties for high- and low-typical members of seventeen semantic categories and found that category members rated low in typicality had fewer properties listed than members high in typicality. Furthermore, the mean number of properties generated was both more highly correlated with typicality ratings than other measures including property overlap with the

superordinate, and a better predictor of reaction time in a property verification task than rated typicality (Ashcraft, 1978a). Ashcraft concluded that the larger number of properties listed for typical than atypical members reflects the fact that subjects are more familiar with the typical members and hence know and can produce more information about them than they can about the less typical members. He further suggested that the mean number of properties produced for an exemplar, as a measure of the amount of readily accessible information, is the variable underlying standard typicality effects in the semantic memory literature.

Ashcraft's data provide the best available support for a familiarity account of typicality ratings. And if familiarity can explain typicality ratings, then the pattern of distribution of properties discussed by Rosch and Mervis (1975), and others, may be artifactual. Thus, those birds, trees, vehicles, etc. that are familiar (and hence typical) are most likely familiar because they are frequent, and frequent because they have properties that made them particularly useful or well-suited to the environment. The properties that make one member of a category well-suited to a given environment will tend to be the same properties necessary for another member to be well-suited, and hence the most frequent or familiar members will have a high degree of property overlap with one another. The less well-

suited members, on the other hand, may have any number of different properties that lead to their lesser suitability. Rosch and Mervis' (1975) family resemblance distribution of properties is thus obtained. Greater familiarity with typical items can also easily be made to account for typicality effects on verification times and order of learning. Verification times may be speeded by the existence of many well-integrated pathways in the representation of the concept, as Ashcraft (1978a) suggests, and children are more likely to be taught, and to have more practice with, the names of common, familiar objects.

Even if familiarity is an important determinant of typicality, can it account for all the variance in typicality ratings? Many of the category members rated as atypical in the Rosch (1975) norms do appear to be low in familiarity relative to the more typical members. However, certain items, such as chickens, pumpkins, and coconuts, that are rated low in typicality (for the categories birds, vegetables, and fruit, respectively) appear to be as frequent and familiar as many of the more typical category members. These instances suggest that some factor in addition to or other than familiarity is accounting for a certain portion of the variance in typicality ratings.

To get at these issues, we used Ashcraft's (1978b) procedure of free-listing properties to members of

categories, with an emphasis on determining the generality of Ashcraft's results. Ashcraft (1978b) used only three typical and three atypical members of each category, and his sampling procedure is not specified. It is not clear to what extent his data reflect the pattern of familiarity across a wide range of typicality values within a category. Our first study sought to determine whether there is necessarily a positive correlation between typicality and number of properties produced for members of categories. The second study tested whether such a correlation holds for large, random samples of category members. The third study examined a possible confounding between familiarity with the word and familiarity with the word's referent.

Experiment 1

The purpose of this experiment was to determine whether familiarity is necessarily correlated with typicality ratings for a large sample of category members that span the typicality range. Two categories were used for which exemplars had been chosen that were deemed by the experimenters to be at least somewhat familiar to college-student subjects. Rated typicality was then correlated with the mean number of properties produced to each category member. If the main determinant of typicality is the item's familiarity, then rated typicality and mean number of properties produced should correlate positively even when

the sample has been chosen such that no items are completely unknown to subjects.

Method

Subjects. Twenty Stanford students generated properties for category members and an additional nineteen students rated the category members for typicality.

Materials. The items were 20 members of the category Furniture and 15 members of the category Bird (see Appendix 1). The Furniture members were the sample used by Rosch and Mervis (1975), which had been selected to span the typicality range; the Bird members were taken from Rosch's (1975) norms, so as to span the typicality range yet not be unknown to college students.

Procedure. Subjects were given sheets of paper with the category name typed in capital letters at the top and the members listed in random order below, and were asked to rate each member for how typical an example of the category it was. Two random orders of members were used for each category. Ratings were made on a scale of 1 to 7, with 7 indicating highest typicality. Property lists were collected from a separate group of subjects. Each item was typed at the top of a sheet of paper. Subjects were given 75 seconds to list all the properties they could think of for each item. Each subject was given all 35 category

members (20 Furniture and 15 Bird members randomly intermixed) in a different random order.

Results and Discussion

The total number of properties listed for each item was averaged over subjects to yield a mean for each item. Pearson correlations were then calculated between the mean number of properties produced and mean typicality ratings. For Bird, the correlation was $-.63$, $p < .025$; for Furniture, the correlation was $-.23$, $p > .10$. Contrary to Ashcraft's (1978b) results, then, the number of properties increased as typicality decreased for both of the present categories. These results are not due to the presence of items at the low end of the typicality scale that do not belong to the category they are in but rather might be typical members of some other category: for eliminating bat from Bird and the electrical appliances (stove, telephone, clock, and radio) from Furniture did not substantially change the two correlations, $r = -.61$ and $r = .25$, $p < .025$ and $p > .10$.

Thus atypical members of a category are not necessarily those that people have had little contact with and know little about. Our atypical Furniture members included the items lamp, piano, and stove, which are likely to be no less frequent in the environment than sofas or desks, and for which equally as many properties were listed. In the Bird sample, subjects produced surprisingly few properties for

some items they rated as typical, such as swallows and mockingbirds, which suggests that these items were actually relatively little known to subjects except that they had generally ordinary bird-like properties. Subjects knew much more about some of the items that were low in typicality, such as chicken and penguin.[1]

What factors account for the discrepancy between our results and those of Ashcraft (1978b)? One likely possibility is sampling procedure. Our items were not a random sampling of all possible members of the two categories; rather, we chose the Bird items so that they would be recognizable to college student subjects, and it is likely that the Furniture items had been selected similarly by Rosch and Mervis (1975). It may be that, while familiarity alone does not determine typicality, it does play some role; in a completely random sampling of items that range widely in familiarity, those items about which subjects know the least may tend to be at the lower end of rated typicality. This was tested in Experiment 2.

Experiment 2

Experiment 2 tested whether, for a random sample of fifteen items from each of eight categories, a positive correlation between typicality ratings and number of properties listed would be found.

Method

Subjects. The subjects were 260 Stanford undergraduates who participated either for course credit or as paid volunteers. 240 of them provided property lists, and 20 provided typicality ratings.

Materials. Eight categories were chosen from the seventeen used by Ashcraft (1978b). Four of them--Furniture, Vehicle, Fruit, and Clothing--were chosen to be at the superordinate level of abstraction, and four--Trees, Bird, Fish, and Flowers--were at the basic level (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976).[2] For each category, we chose fifteen members (see Appendix 2). Six of them were the six used by Ashcraft (1978b). An additional nine were sampled from Battig and Montague (1969), primarily using exemplars with a production frequency greater than or equal to nine.

The words were typed at the top of sheets of paper and assembled into packets. Each packet contained one member from each of the eight categories. Packets of items were constructed such that no more than 4 packets out of 240 contained the same set of items, and no more than 2 of those 4 contained the items in the same order.

Procedure. Typicality ratings were collected in the same manner as in Experiment 1. Each subject rated all

eight categories in a different random order. There were two random orderings of each set of fifteen category members, with half the subjects receiving one set of random orderings and half the other.

Subjects providing property lists were each given a packet containing one member from each of the eight categories. They had 75 seconds to list properties for each member. They were explicitly told to write as much as they knew for any item they were uncertain about.

Results and Discussion

Mean typicality ratings and mean number of properties produced were calculated for each item, and the correlation between these two measures was then obtained for each category. In contrast to Experiment 1, the correlation coefficients were positive for all eight categories, indicating that subjects knew more about exemplars high in typicality than about those low in typicality. For five of the categories, the correlations were significant: Birds, r = .68; Flowers, r = .69; Fruit, r = .66; Tree, r = .67; and Vehicle, r = .62, p < .025 in all cases. The remaining three correlations were positive but non-significant: Clothing, r = .24; Fish, r = .07; and Furniture, r = .36, p > .10 in all cases. There was no systematic difference between basic and superordinate level categories. Also, as expected, the degree of familiarity varied much more for

these items than for the set in Experiment 1. Whereas the mean number of properties listed ranged from 5.95 to 9.60 in Experiment 1, the range for Experiment 2 was 1.56 to 11.63. Thus, when the items range from very familiar to very unfamiliar, typicality ratings do show a substantial correlation with familiarity, as Ashcraft (1978b) found.

The change in sampling procedure from Experiment 1 to Experiment 2 reversed the direction of the critical correlations. It is not clear, however, exactly how the greater range of familarity in Experiment 2 items might have changed the direction of the results. The familiarity explanation of typicality assumes that subjects have at least a rough idea of the appearance of the referent of the item they are rating as low in typicality; it is the relative infrequency of that referent in the environment that leads it to be perceived as less typical, and ability to list fewer properties is a by-product and indicator of this lesser familiarity. However, the extreme lack of knowledge evidenced by many subjects for a number of Experiment 2 items suggests an alternative explanation. If subjects have no idea what the referent of a word is, they cannot make a typicality judgement about the referent itself and perhaps resort to a strategy of assigning low ratings to such words. Experiment 3 was designed to distinguish between the two possibilities.

Experiment 3

In most typicality-rating tasks, subjects are given a scale that forces them to place every item on the typicality continuum. Subjects are never given a chance to indicate items which they are completely unfamiliar with. Experiment 3 collected a new set of typicality ratings in which an alternative response to a typicality number was U, indicating the rater was too unfamiliar with the referent of the word to rank it on the scale. If some of the low-rated items in Experiment 2 were such that subjects had no idea of the referents' appearances, then these items should receive U ratings from the present group of subjects. Furthermore, when items rated as U by a substantial number of subjects are removed from the calculations, the correlations between number of properties and typicality should be reduced. On the other hand, if the typicality ratings in Experiment 2 truly reflect perceived typicality of the referent, then the number of U ratings should be minimal and there should be no tendency for them to cluster at the lower end of the typicality range.

Method

Subjects. Twenty Stanford undergraduates participated in partial fulfillment of a course requirement.

Materials. Packets of eight rating sheets, identical to those in Experiment 2, were used.

Procedure. Subjects were given standard typicality-rating instructions, and were further instructed that, "For any item which is unfamiliar enough to you that you don't feel able to accurately rank it with respect to the others, place a U in the blank instead of a number."

Results and Discussion

The number of times each item was given a U rating was tabulated. Out of 120 items, 30 were rated as U by at least one subject. To determine whether U ratings clustered at the lower end of the typicality range, each category was divided into the upper and lower halves of the typicality range (as determined by the ratings in Experiment 2), and the number of U ratings in each half of the range was counted for each category and then averaged across the eight categories. For the sixteen subjects who gave U ratings, all gave more to category members in the lower half of the typicality range. This difference was significant by a Wilcoxon matched-pairs signed-ranks test, $p < .01$.

It might also be noted that for three of the categories--Clothing, Furniture, and Vehicle--there were virtually no U ratings given by any subject, while for the remaining categories as many as seven members received Us from a single subject. These three categories were all at the superordinate level, and hence the names of their members should be at the basic or most commonly used level

of abstraction (Rosch et al., 1976). Two of these categories, Clothing and Furniture, were two of the three categories in Experiment 2 for which the correlation between typicality and number of properties was not significant, which is consistent with the relation between typicality and familiarity being largely due to unfamiliar words.

The Experiment 2 correlations between typicality and mean number of properties were recalculated, omitting all items that received a U rating from four or more subjects in Experiment 3 (one fifth of the number who provided ratings). Table 1 gives the correlations before and after omitting the U items. Correlations dropped for all five categories that contained members with the criterial number of U ratings. One of these categories, Fish, which was close to zero previously, became slightly negative. Two categories, although remaining positive, dropped below the level of significance. For the remaining two categories, the correlations remained significant at the .05 level, though not at the .01 level. (Lowering the criterion for omitting exemplars from four to three U ratings caused the correlation for Fruit to drop to a nonsignificant .49, and also caused the correlation for Fish to become more negative, -.44.)

Thus, after omitting words for which subjects did not have minimal knowledge of the referent, only two out of

Table 1

Correctional Between Rated Typicality and Number of Properties Produced, Before and After Removal of Unfamiliar (U) Exemplars.

	Before Removal (Experiment 2, n = 15 for each)	After Removal
Birds	.68**	.60* n = 13
Clothing	.24	.24 n = 15
Fish	.07	-.06 n = 13
Flowers	.69**	.62 n = 9
Fruit	.66**	.62* n = 14
Furniture	.36	.36 n = 15
Tree	.67**	.44 n = 13
Vehicle	.62*	.62* n = 15

* = $p < .05$, two-tailed

** = $p < .01$, two-tailed

eight correlations remained statistically significant. All except one did, however, remain positive. This suggests there may be some influence of familiarity of referents on perceived typicality, as claimed by Ashcraft (1978b).

General Discussion

These experiments indicate that familiarity is not the major determinant of typicality. Experiment 1 showed that subjects are not necessarily less familiar with items they rate low in typicality than those high in typicality, where familiarity was measured by number of properties produced. Experiment 2 showed that in a large, random sample of items, there was a tendency for atypical items to have few properties listed, replicating Ashcraft (1978b), but Experiment 3 suggested that part of the reason for this result may be a response bias toward assigning totally unfamiliar items to the bottom of the scale, rather than reflecting low perceived typicality of the referents themselves. These results undermine Ashcraft's (1978a) suggestion that familiarity is the variable underlying typicality effects in semantic memory tasks such as verification times. (In Experiment 1, familiarity and typicality were negatively correlated; so if mean number of properties produced is indeed a better predictor of reaction times than rated typicality, as Ashcraft found, then reaction time should be faster to items at the lower end of

the Birds sample used in Experiment 1, a prediction that contradicts a substantial body of findings, e.g., Smith et al., 1974).

In addition to the present results, there is another finding with natural categories indicating that typicality differences need not arise from variations in familiarity. Smith et al. (1974) found that the relative typicality of two exemplars could change depending on what category was being considered. For example, while a robin was rated as a more typical bird than a chicken, chickens were considered to be the more typical animal. It would seem to be true for many items that their typicality may vary depending on the category they are rated in relation to. A snake may be a typical reptile, a moderately typical vertebrate, and an atypical animal.

It therefore seems that differences in typicality can exist independent of any level of familiarity. The fact that most of the correlations in Experiment 3 remained positive after removal of totally unfamiliar items suggests, however, that familiarity played some part in determining the typicality ratings of Experiment 2 items. Category members that are recognizable to subjects but still lower in familiarity than others may tend to be rated lower in typicality than the familiar ones, particularly if their names are not frequently heard. Perhaps the safest

conclusion is that more than one factor can influence typicality ratings: Category members that are familiar may be low in typicality if, like chickens, coconuts, or subways, they have properties uncommon within the category; others, such as dogwood trees, ravens, and cherry blossoms, may be given low ratings if they are less familiar to subjects than many of the other category members; and some, such as catbirds, jonquils, and ginko trees, may be assigned low ratings despite their similarity to typical members of their categories because subjects do not know their appearance. The extent to which the various factors come into play may depend on the general level of familiarity of the exemplars. For categories such as Clothing, where most exemplars are quite well-known, property-similarity may be the main determinant of ratings, while for other, such as Trees, familiarity may be more heavily weighted.

Appendix 1

Items for Experiment 1,
in descending order of typicality

Birds	Furniture
robin	sofa
bluebird	chair
seagull	table
swallow	desk
falcon	dresser
mockingbird	bed
starling	bookcase
owl	piano
vulture	footstool
sandpiper	lamp
chicken	mirror
flamingo	cushion
albatross	vase
penguin	clock
bat	rug
	picture
	radio
	stove
	closet
	telephone

Appendix 2

Items for Experiments 2 and 3,
in descending order of typicality

Birds	Clothing	Fish	Flowers
sparrow	shirt	trout	rose
robin	dress	salmon	daisy
bluejay	slacks	tuna	poppy
crow	coat	goldfish	lily
hawk	socks	minnow	iris
wren	underpants	carp	marigold
duck	belt	sardine	African violet
owl	sweatshirt	shark	lilac
mockingbird	bathrobe	whitefish	azalea
chicken	scarf	sunfish	cherry blossom
raven	gloves	shrimp	peony
thrush	watch	dolphin	begonia
pelican	necklace	lobster	jonquil
catbird	cape	walleye fish	sweet pea
albatross	cane	eel	dogwood blossom

Fruit	Furniture	Tree	Vehicle
apple	chair	pine	car
peach	table	oak	truck
strawberry	sofa	maple	airplane
pear	bureau	redwood	bicycle
grape	lounge chair	elm	train
blueberry	television	sequoia	van
lemon	bench	palm tree	jeep
watermelon	shelf	beech tree	subway
raisin	rug	peach tree	cable car
fig	mirror	pear tree	feet
coconut	cushion	cypress	rowboat
pumpkin	chaise lounge	dogwood	horse
pawpaw	cedar chest	mimosa	raft
olive	drapes	ginko	go-cart
gourd	vase	bamboo	dogsled

Footnotes

[1] There is a distinction to be made between category members such as robins that are familiar because they are probably frequently encountered in the environment, and those such as penguins for which the knowledge of the many properties listed is more likely to be acquired from stories and pictures. The property-listing method does not distinguish between these two different kinds of familiarity. But it is doubtful that our results in any way hinge on this distinction, since all the Furniture members, atypical as well as typical, are probably equally likely to be encountered in the environment.

[2] We favored Rosch et al's (1976) biological basic-level categories because they tended to have more members with relatively simple names.

References

Ashcraft, M. H. Property dominance and typicality effects in property statement verification. Journal of Verbal Learning and Verbal Behavior, 1978a, 17, 155-164.

Ashcraft, M. H. Property norms for typical and atypical items from 17 categories: A description and discussion. Memory and Cognition, 1978b, 6, 227-232.

Battig, W. F., & Montague, W. E. Category norms for verbal items in 56 categories: A replication and extension of the Connecticut category norms. Journal of Experimental Psychology Monograph, 1969, 80 (3, Part 2).

Kucera, H., & Francis, W. N. Computational Analysis of Present-Day American English. Providence, R.I.: Brown University Press, 1967.

McCloskey, M. The stimulus familiarity problem in semantic memory research. Journal of Verbal Learning and Verbal Behavior, 1980, 19, 485-502.

Mervis, C. B., Catlin, J., & Rosch, E. Relationships among goodness-of-example, category norms, and word frequency. Bulletin of Psychonomic Society, 1976, 7, 283-284.

Mervis, C. B., & Rosch, E. Categorization of natural objects. Annual Review of Psychology, in press.

Rips, L. J., Shoben, E. J., & Smith, E. E. Semantic distance and the verification of semantic relations. Journal of Verbal Learning and Verbal Behavior, 1973, 12, 1-20.

Rosch, E. Cognitive representations of semantic categories. Journal of Experimental Psychology: General, 1975, 104, 192-233.

Rosch, E., & Mervis, C. B. Family resemblances: Studies in the internal structure of categories. Cognitive Psychology, 1975, 7, 573-605.

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. Basic objects in natural categories. Cognitive Psychology, 1976, 8, 382-439.

Rosch, E., Simpson, C., & Miller, S. Structural bases of typicality effects. Journal of Experimental Psychology: Human Perception and Performance, 1976, 2, 491-502.

Smith, E. E., Shoben, E. J., & Rips, L. J. Structure and

process in semantic memory: A featural model for semantic decisions. Psychological Review, 1974, 81, 214-241.